

# CS354: Machine Organization and Programming

Lecture 26  
Monday the November 02<sup>nd</sup> 2015

Section 2  
Instructor: Leo Arulraj

© 2015 Karen Smoler Miller  
© Some examples, diagrams from the CSAPP text by Bryant and O'Hallaron

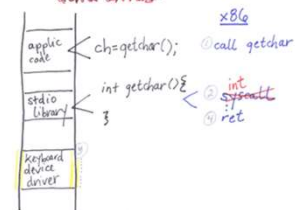
## Class Announcements

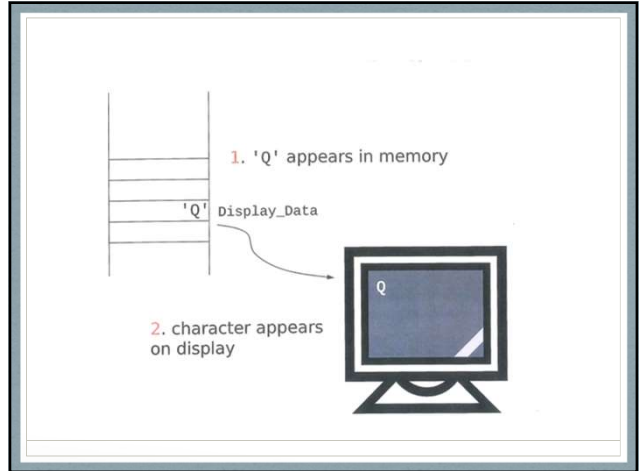
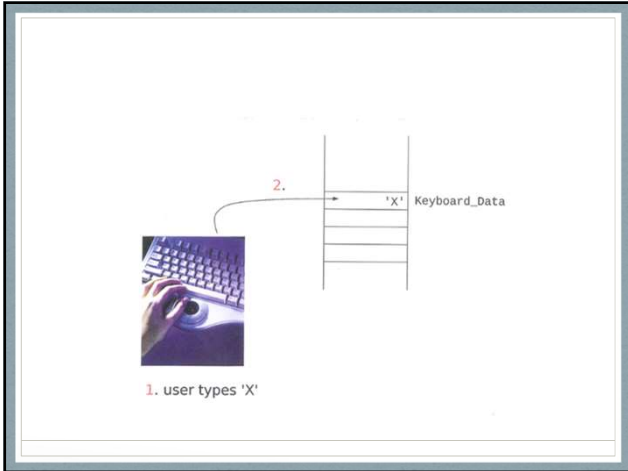
1. Programming Assignment 3 is due by 9 AM day after tomorrow (11/04 – Wednesday). As usual, you can submit it upto 48 hours after the deadline with penalties.

## Lecture Overview

1. Interrupts and Exception
2. Intro to Processes

To control access to individual devices, we place code that communicates with the device into the OS. (this portion of the OS is also called the *kernel*) and in special routines called *device drivers*.





The Keyboard driver code:

```

movl Keyboard_Data, %eax
ret from syscall
    
```

↑ Assume this is where the ASCII char is supposed to be after syscall completes.

The display driver code:

```

movl %eax, Display_Data
ret from syscall
    
```

What happens if user has not typed a char on the keyboard?

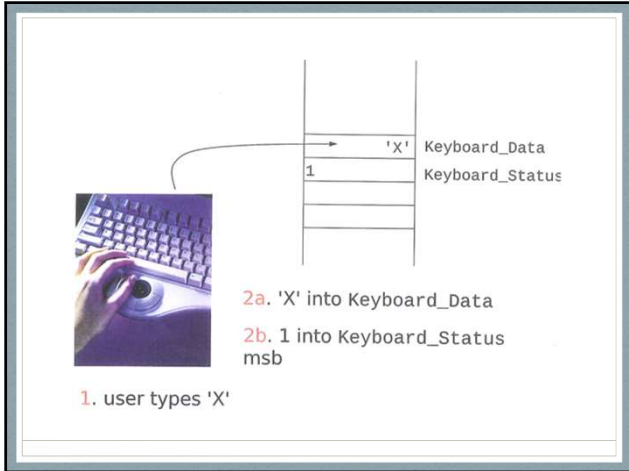
We want **blocking input**.  
(no ret from syscall until there is a char)

We need a **status bit**

1 ready  
0 not ready (busy)

Place this bit into its own memory mapped word, or make it the msb, so code can test for  $\geq 0$  or  $\leq 0$ .

Keyboard\_Data  
Keyboard\_Status  
Display\_Data  
Display\_Status

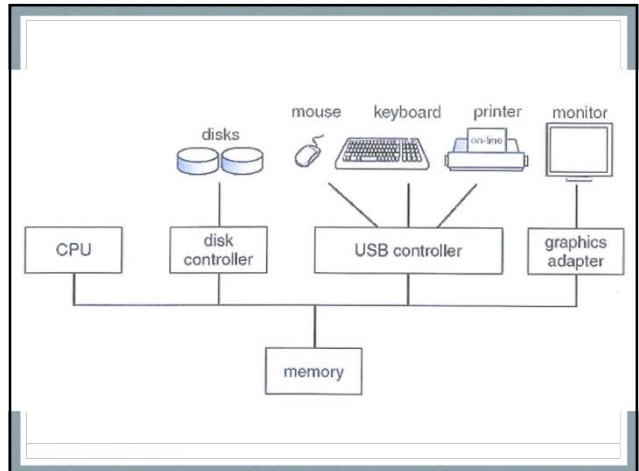
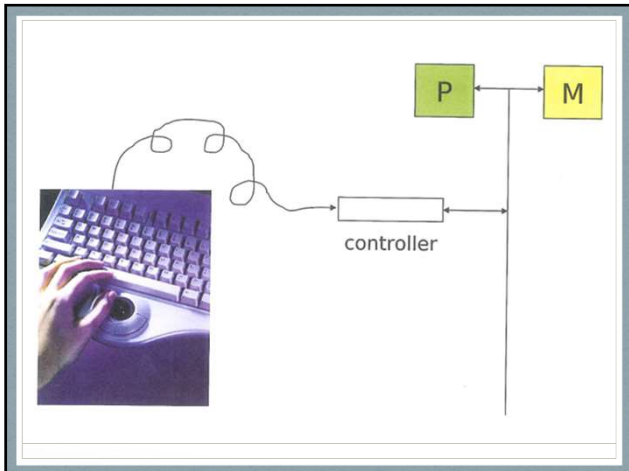


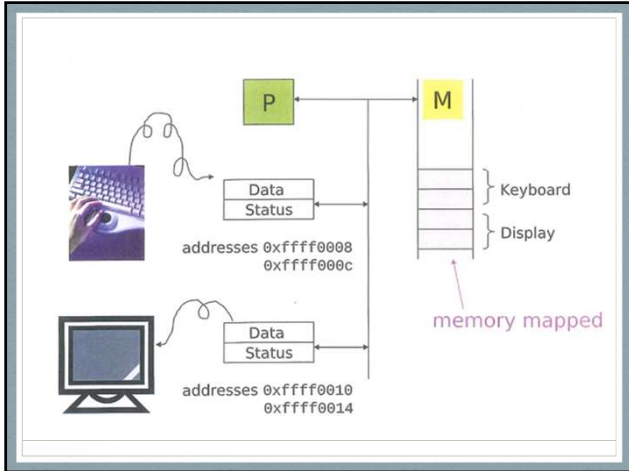
Now, driver code uses a **Spin wait loop** (to implement blocking I/O)

```

Kb_spin: testl Keyboard_Status, Keyboard_Status
        jz   Kb_spin
        movl Keyboard_Data, %eax
        ret from %eax

disp_spin: testl Display_Status, Display_Status
          jz   disp_spin
          movl %eax, Display_Data
          ret from %eax
    
```





Byte transfers are OK,  
But, what about faster devices that like to transfer more than a byte ?

the solution: **DMA**  
Direct Memory Access

**Issue for spin wait loop implementations:**

One byte *only* in `_Data` has the potential for an incorrect result.

For example, if the user types 2 characters on the keyboard before `getchar()` is called.

The needed fix introduces a kernel-maintained **queue** for each device.

Then, the kernel **polls** to check status bits and handle any ready devices.

INTERRUPTS

here is an analogy. . .

Teacher  $\longleftrightarrow$  is  $\longleftrightarrow$  OS

each student  $\longleftrightarrow$  is  $\longleftrightarrow$  I/O device

Consider the *inefficiency* of OS polling.

1

©These slides may be freely used, distributed, and incorporated into other works.

**I**  
**N**  
**T**  
**E**  
**R**  
**R**  
**U**  
**P**  
**T**  
**S**

Because polling is *so inefficient*,

instead of

```
graph TD; OS[OS] -- ask --> D1[device 1 ready?]; OS --> D2[device 2 ready?]; OS --> Dn[device n ready?];
```

2

©These slides may be freely used, distributed, and incorporated into other works.

**I**  
**N**  
**T**  
**E**  
**R**  
**R**  
**U**  
**P**  
**T**  
**S**

Turn the situation upside down

```
graph TD; D1[device 1] -- ready! --> OS[OS]; D2[device 2] --> OS; Dn[device n] --> OS;
```

3

©These slides may be freely used, distributed, and incorporated into other works.

**I**  
**N**  
**T**  
**E**  
**R**  
**R**  
**U**  
**P**  
**T**  
**S**

©These slides may be freely used, distributed, and incorporated into other works.

### Anatomy of an exception

An exceptions is an abrupt change in control flow.

Examples: div by 0, arithmetic overflow, page fault, I/O request completes, Ctrl-C

```
graph LR; subgraph Application_program; direction TB; A[Application program]; end; subgraph Exception_handler; direction TB; H[Exception handler]; end; Event[Event occurs here] -.-> A; A -- "curr / next" --> H; H -- "Exception" --> A; H -- "Exception return (optional)" --> A; A --> A; H --> H; A --> H; H --> A;
```

©These slides may be freely used, distributed, and incorporated into other works.


## Classes of Exceptions

1. **Interrupts (Asynchronous):** Always return to next instruction.
2. **Traps & System Calls (Synchronous):** Always return to next instruction.
3. **Faults (Synchronous):** Might return to next instruction.
4. **Aborts (Synchronous):** Never returns

## Interrupts

Examples of Interrupts:

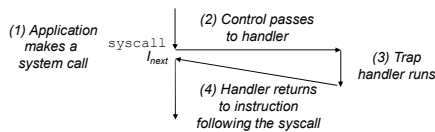
- Timer interrupt
- Arrival of a packet from a network
- When a key is pressed on the keyboard
- When the mouse is moved



## Traps

Traps are intentionally issued by executing an instruction.


Example: System calls



## Faults

Faults result from error conditions that might be correctable.


Examples: Page fault, Divide error



## Aborts


Aborts result from unrecoverable fatal errors.

Example: parity errors due to DRAM bit corruption



## Exception Table


Exception table



## Exception Table lookup

Exception is similar to procedure calls except for some important differences:

- Return address is not the next instruction always
- Push EFLAGS register also onto kernel stack
- Run exception handler in kernel mode



## IA32 Exception Table

From CSAPP text book:

Exception Number	Description	Exception Class
0	Divide error	Fault
13	General protection fault	Fault
14	Page fault	Fault
18	Machine check	Abort
32-127	OS-defined	Interrupt or trap
128 (0x80)	System call	Trap
129-255	OS-defined	Interrupt or trap

Rather important, but not covered in textbook:

If running a handler, and a new interrupt request arrives, what should happen?


\* Continue on, complete handling of current interrupt, then, when done, deal with new request?  
(probably) nonreentrant

\* Interrupt the handling of this interrupt?  
reentrant

(16)

Every architecture has a control bit which identifies whether the P+E cycle is paying attention to IRQs. Called *interrupt enable*.

on x86:  
EFLAGS



fetch+execute cycle:  
1 enabled  
0 disabled

① if IF=1 and interrupt requested, go handle it

- ① fetch instr
- ② PC update
- ③ decode
- ⋮

Consider the x86 instruction:


`cli` clear IF

What happens if an application includes this `cli` instruction?

Irrelevant (to this discussion) x86 instruction:

`sti` set IF

OS relies on clock interrupts to allocate processing time . . .



As the clock interrupts, the kernel runs, and it decides which program runs next.



Clarified instruction:  
`cli` clear IF if CPL is high enough,  
 otherwise trap

Does `sti` also need to be a privileged instruction?

Keep  $IF=1$  while  $CPL=00$ .  
 (So, applications can always be interrupted)

HW must disable interrupts while saving state + at least until first instruction within handler is fetched.

Better Definitions:  
 nonreentrant  $IF=0$  the entire time a handler runs.  
 reentrant interrupts may be reenabled while handler runs (usually only for higher priority requests)

